arXiv:1601.07752v2 [math.NA] 1 Feb 2016

Enhancing the Power of Cardinal's Algorithm

Victor Y. Pan

Department of Mathematics and Computer Science
Lehman College of the City University of New York
Bronx, NY 10468 USA
Ph.D. Programs in Mathematics and Computer Science
The Graduate Center of the City University of New York
New York, NY 10036 USA
victor.pan@lehman.cuny.edu
http://comet.lehman.cuny.edu/vpan/

Abstract

Cardinal's factorization algorithm of 1996 splits a univariate polynomial into two factors with root sets separated by the imaginary axis, which is an important goal itself and a basic step toward root-finding. The novelty of the algorithm and its potential power have been well recognized by experts immediately, but by 2016, that is, two decades later, its practical value still remains nil, particularly because of the high computational cost of performing its final stage by means of computing approximate greatest common divisor of two polynomials. We briefly recall Cardinal's algorithm and its difficulties, amend it based on some works performed since 1996, extend its power to splitting out factors of a more general class, and reduce the final stage of the algorithm to quite manageable computations with structured matrices. Some of our techniques can be of independent interest for matrix computations.

2000 Math. Subject Classification: 65H05, 65F15, 30C15, 26C10, 12Y05

KEYWORDS: Polynomials, Factorization, Root-finding, Companion matrix, Frobenius algebra

1 Introduction

The solution of a polynomial equation or equivalently root-finding for a univariate polynomialhas been a central problem in Math and Computational Math for four millennia [B40], [B68] and is still important for robotic and control, geometric modeling, computer algebra, signal and image processing, computer vision, tensor decomposition, and financial mathematics (see, e.g., [MP13, Preface]). Moreover root-finding is closely linked to univariate polynomial factorization, which has important applications to the time series analysis, Weiner filtering, noise variance estimation, covariance matrix computation, and the study of multi-channel systems (see [W69], [BJ76], [B83], [DM89], [DM90], [VD94]). Hundreds (if not thousands) efficient techniques and algorithms a have been proposed for these tasks, and so the designers of new univariate polynomial root-finders are naturally prompted to draw from a vast reservoir of efficient algorithms and techniques of their predecessors. Some recent papers, e.g., [PT13] and [BSSY], combine such techniques and algorithms, complement them with a few more recent methods, and arrive at nearly optimal root-finders in terms of their Boolean complexity. Can we produce new efficient algorithms with practical promise in this way? Although the Boolean complexity is a fundamental concept and properly measures the efficiency of polynomial root-finders, other criteria can be also important. E.g., the current best

package of root-finding subroutines MPSolve is not supported by competitive Boolean complexity estimates, but nevertheless is loved by the user.

Presently we enhance an advanced algorithm of [C96] by adding novel numerical and symbolic techniques for computations with general and structured matrices, randomization, and derandomization (see Sections 5–8). Originality and potential power of that algorithm have been appreciated by experts immediately, in 1996, but two decades later its practical value still remains nil, particularly because of high computational cost of its final stage of computing an approximate greatest common divisor of two polynomials. (Hereafter we use the acronym *AGCD*.) We briefly recall Cardinal's algorithm and its difficulties, amend it based on some works performed since 1996, extend its power to splitting out all its factors whose roots lie in a family of boxes that tile the complex plane, and replace the AGCD computation by the approximation of eigenspaces of the associated companion matrix. These techniques decrease the arithmetic cost, stabilizes the computations numerically, and can be of independent interest for matrix and polynomial computations.

2 Some basic definitions

Hereafter we write "*Gaussian*" for "standard Gaussian random". In the Appendix we recall some definitions and basic results for computations with polynomials, structured matrices and Gaussian matrices and for eigen-decomposition of a matrix. Next we list some basic definitions for matrix computations (cf. [GL13]).

1. I_g is a $g \times g$ identity matrix. $O_{k,l}$ is the $k \times l$ matrix filled with zeros.
2. $(B_1 \mid B_2 \mid \cdots \mid B_k)$ denotes a block vector of length k with blocks $B_1, B_2, \dots, B_k$.
   $\mathrm{diag}(B_1, B_2, \dots, B_k)$ denotes a $k \times k$ block diagonal matrix with blocks $B_1, B_2, \dots, B_k$.
3. M^T denotes the transpose of a matrix M. M^H denotes its Hermitian transpose.
4. An $m \times n$ matrix M is called *unitary* if $M^H M = I_n$ or $MM^H = I_m$.
5. Any matrix M of rank ρ has *compact SVD*, $M = S_{M,\rho}\Sigma_{M,\rho}T^H_{M,\rho}$, (hereafter referred to just as *SVD*), where $S_{M,\rho}$ and $T_{M,\rho}$ are unitary matrices of singular vectors and $\Sigma_{M,\rho} = \mathrm{diag}(\sigma_j(M))_{j=1}^{\rho}$ is the diagonal matrix of singular values, $\sigma_1(M) \ge \sigma_2(M) \ge \cdots \ge \sigma_\rho(M) > 0$. ($||U|| = ||U^+|| = 1$, $||UM|| = ||M||$ and $||MU|| = ||M||$ if U is a unitary matrix.)
6. M^+ denotes the *Moore–Penrose pseudo inverse* of the matrix M. ($M^+ = T_{M,\rho}\Sigma^{-1}_{M,\rho}S^H_{M,\rho}$ if $M = S_{M,\rho}\Sigma_{M,\rho}T^H_{M,\rho}$. $||M^+|| = 1/\sigma_\rho(M)$. $M^+ = M^{-1}$ for a nonsingular matrix M.)
7. $\kappa(M) = ||M||\,||M^+|| = \sigma_1(M)/\sigma_\rho(M) \ge 1$ denotes the *condition number* of a matrix M. A matrix W is *ill-conditioned* if it is close to a matrix of a smaller rank or equivalently if its condition number is large in context. The matrix is *well-conditioned* if its condition number is reasonably bounded.

3 Cardinal's factorization algorithm and its extensions

In his paper [C96], Cardinal assumes that he is given a univariate polynomial of degree n with real or complex coefficients and with no roots x_j lying on the imaginary axis $\{x:\ \Re x = 0\}$,

$$p(x) = \sum_{i=0}^{n} p_i x^i = p_n \prod_{j=1}^{n}(x - x_j), \quad p_n \prod_{j=1}^{n} \Re x_j \neq 0. \tag{3.1}$$

Then he fixes the initial polynomial $f_0(x) = x$ and applies recursively the Möbius' map

$$f_{i+1}(x) \to \frac{1}{2}\Big(f_i(x) + \frac{1}{f_i(x)}\Big) \mod p(x), \tag{3.2}$$

for $i = 0, 1, \dots$ until convergence.

In the following estimates we write (cf. [H08, page 107])

$$\text{sign}(x) = \text{sign}(\Re x) \text{ for all complex numbers } x.$$

Theorem 3.1. *(Cf. [C96, Lemma 1], [BP96, page 500].)*
Fix a complex $x = x^{(0)}$ and define iterations (3.2). Then

$$|x^{(i)} - \text{sign}(x)| \le \frac{2\tau^{2^i}}{1-\tau^{2^i}} \text{ for } \tau = \Big|\frac{x - \text{sign}(x)}{x + \text{sign}(x)}\Big|, i = 0, 1, \dots$$

The theorem shows global quadratic convergence of the iterations, right from the start.

As soon as the iterates $x^{(i)}$ approach $\text{sign}(x)$ closely, Cardinal approximates the two factors of the polynomial $p(x)$ whose root sets are separated by the imaginary axis.

Remark 3.1. *We can map the axis $\Im$ into the unit circle $C(0,1) = \{x : |x| = 1\}$ (cf. Theorem 11.3), then estimate the distances of all roots from the origin (cf. [S82, Section 14, Task S]), map the circle $C(0,1)$ back into the axis $\Im$, obtain some estimates for the distances of the roots from the axis and hence for the numbers of iterations (3.2) required for convergence. The algorithm of [S82] has a reasonable Boolean cost, but is not good numerically, but one can replace it with the algorithm implied by [B96, Theorems 5 and 6], which is much simpler, numerically stable, and produces desired estimates for a large input class.*

Hereafter "*op*" stands for "arithmetic operation".

Remark 3.2. *The real parts $\rho_j = \Re(x_j)$ of the roots of $p(x)$ on the axis $\Im$ satisfy the two equations $\Re(p(\rho_j)) = 0$ and $\Im(p(\rho_j)) = 0$. For most inputs we can check rationally (by using only ops) if the axis contains any root. Similarly, for most input polynomials $p(x)$ with nonreal coefficients, we can check if they have real roots, which is non-rational task if all coefficients are real. We can compute the numbers of roots in a disc and hence on each side of the axis $\Im$, by applying the algorithms of [R87, Section 7], [BP00], or [BSSY].*

By applying the maps $x \to y = ax + b$ and $x \to y = \frac{x-r}{x+r} + b$ for appropriate real $r \neq 0$ and complex a and b, $|a| = 1$, we can move any straight line L as well as any circle C on the complex plane into the imaginary axis (cf. Theorems 11.1 and 11.3). If no roots of the resulting polynomial $q(y)$ lie on the imaginary axis, then the algorithm splits that polynomial about the imaginary axis, and this is readily translated into splitting the polynomial $p(x)$ into two factors whose root sets are separated by the line L or the circle C.

By applying such a factorization k times, we can split out a factor of the polynomial $p(x)$ with the root set in a region bounded by k straight lines and circles on the complex plane. For $k = 2$ this can be a rectangular band, annulus, angle, or half-disc; for $k = 3$, say, a triangle or a sector in a disc; for $k = 4$ a square or rectangle, etc. We can cover all roots of a polynomial $p(x)$ with a square having vertical and horizontal sides (cf. Theorem 11.4), partition it by vertical and horizontal lines into rectangles (*tiles*), and split out the factors of polynomial $p(x)$ having root sets in selected rectangles (see Section 9). We can apply the algorithm concurrently over all lines and all selected rectangles with minimum synchronization and communication among processors.

4 Operations in Frobenius algebra and their complexity

Every Möbius' iteration of equation (3.2) amounts to computing a reciprocal of a polynomial modulo $p(x)$. We can perform that operation in $O(n \log^2(n))$ by applying numerically unstable algorithms in [P01, Algorithm 2.7.9], [P01, Section 2.10 and Chapter 5]) lops, or the alternative solution implicit in [C96], which reduces root-finding for a polynomial $p(x)$ to the approximation of the eigenvalues

of its $n \times n$ companion matrix

$$C_p = \begin{pmatrix} 0 & & & & -p_0/p_n \\ 1 & \ddots & & & -p_1/p_n \\ & \ddots & \ddots & & \vdots \\ & & \ddots & 0 & -p_{n-2}/p_n \\ & & & 1 & -p_{n-1}/p_n \end{pmatrix};$$

they are precisely the roots of $p(x)$. The *Frobenius algebra* $\mathcal{A}_p$ of polynomials in the matrix C_p is isomorphic to the algebra of polynomials in a variable x modulo $p(x)$. The matrices of this algebra have Toeplitz-like structure and have displacement rank at most 2 (see [C96, Proposition 2] for a proof and our Section 14 for definitions).

Performing the ith iteration (3.2) can be reduced to solving a Toeplitz-like linear system of equations with the coefficient matrix $f_i(C_p)$. For this [C96] proposes to apply MBA [M80] and [BA80], which also uses $O(n \log^2(n))$ flops and is also prone to problems of numerical stability [B85]. By modifying [C96], [BP96] computes fewer reciprocals modulo $p(x)$, but now we can perform this operation more efficiently. Let us summarize the state of the art in this area, including acceleration of the MBA algorithm by a factor of $\log(n)$ and its numerical stabilization.

Theorem 4.1. *Numerically stable computations involve $O(n)$ flops for an addition and $O(n \log(n))$ for multiplication and inversion in the Frobenius algebra $\mathcal{A}_p$ as well as for the solution of a linear system of equations with a matrix in that algebra.*

Proof. Addition in $\mathcal{A}_p$ is trivial. The algorithms of [C96] and [P05] support the cost bound for multiplication and ones of [XXG12], [XXCB14], and [P15, Section 9.8] for inversion. □

5 Heuristic transformations of companion matrix and its polynomial

One can simplify various computations with a companion matrix C_p by replacing it with a matrix C_q for one of the polynomials $q(y)$ of Theorems 11.1–11.3. Besides the shift and invert technique, well-known for polynomial root-finding (cf. [JT70], [JT70]), but one can include scaling of the variable x and in some applications Dandelin's root-squaring map of Theorem 11.2.

One can further replace C_q by a matrix $C_q + sI_n$ or $(C_q + sI_n)^{-1}$ for a proper real or complex shift value s and apply a nested recursive process combining such maps. The resulting transforms of the matrix C_p can facilitate eigen-solving in various ways. The following three examples demonstrate such potential benefits and can motivate further exploration of the approach.

Example 5.1. *Quite typically, polynomials $p(x)$ of large degree have roots more or less uniformly distributed about a circle centered in the origin (cf. [A66], [HN08], and [DGS14]).*

Consider a sample polynomial $p(x) = x^n - f^n$ where $f \neq 0$.

If $|f|$ is small, then the matrix C_p is close to the derogatory companion matrix C_q of the polynomial $q(x) = x^n$. Approximation of the eigenvalues and eigenspaces of such matrices is an ill-conditioned task, that is, always requires computations with a high precision [S01], [GL13].

In our case, however, we can let $x = fy$ and obtain $p(x) = q(y) = f^n(y^n - 1)$.

Then $C_q = (\omega^{ij})_{i,j=0}^{n-1}$, $\omega = \exp(\frac{2\pi}{n}\sqrt{-1})$, is the matrix of discrete Fourier transform and is a unitary matrix up to scaling by $\sqrt{n}$.

Example 5.2. *Suppose that a matrix C_p is ill-conditioned (that is, nearly singular). Then numerical problems would arise at all Möbius' iterations (3.2) of Cardinal's algorithm, but may disappear if we replace C_p by its transform such as $C_q + cI_n$, $(C_q + cI_n)^{-1}$, $C_{q\mathrm{rev}} + cI_n$, or $(C_{q\mathrm{rev}} + cI_n)^{-1}$ for a proper polynomial $q(x)$ of Theorems 11.1–11.3 and a fixed or random constant c.*

For applications to Cardinal's algorithm, note that inversions, scaling by a real nonzero constant, and shifts by purely imaginary values keep the imaginary axis invariant.

Example 5.3. *Let us increase the relative distances among the roots of a polynomial by combining shift-and-invert maps with lifting root-squaring iterations of Theorem 11.2 and the converse descending technique of [P95], [P02].*

1. *Shift the variable:* $p(x) \to q(y) = p(x)$ *for* $y = x - f$ *and a sufficiently large* $|f|$.
2. *Reverse a polynomial:* $q(y) \to r(z) = q_{\rm rev}(y) = y^n q(1/y)$.
3. *Shift the variable again:* $s(w) \to r(z)$ *for* $w = z + 1/f$.
4. *Apply one or two root-squaring iterations of Theorem 11.2, producing a polynomial* $t(v)$.
5. *Reverse a polynomial again:* $t(v) \to t_{\rm rev}(v)$.
6. *Approximate the roots of the polynomial* $t_{\rm rev}(v)$.
7. *Recover the roots of* $p(x)$ *by combining Theorems 11.1–11.3 and the descending techniques of [P95], [P02].*

6 The final stage of Cardinal's algorithm and our alternative

At the final stage of splitting a polynomial $p(x)$ into the product of two factors in [C96], Cardinal approximates each of them as an approximate greatest common divisor (AGCD) of $p(x)$ and an auxiliary polynomial $q(x)$ of degree n. All known algorithms for such computation are not faster than root-finding for $p(x)$ or $q(x)$ (cf. [P98/01], [BB10], and the references therein). Our approach at this stage is radically different from Cardinal's: we compute the two factors from the two eigenspaces of the matrix C_p associated with them. Iteration (3.2) takes the following equivalent form,

$$X_0 = C_p,\ X_{i+1} = 0.5(X_i + X_i^{-1}),\ i = 0, 1, \dots, \tag{6.1}$$

and Theorem 3.1 implies global quadratic convergence of the eigenvalues of the matrices X_i to ± 1. Hence the eigenvalues of the matrices $X_i \pm I_n$ converge to -2 and 0 and to 0 and 2, respectively. We call the absolutely small eigenvalues and the associated eigenspaces *dominated* and call the remaining eigenvalues and the associated eigenspaces *dominant*. The inversion of a matrix swaps its dominant and its dominated eigenvalues and eigenspaces.

The following algorithm reduces the splitting of the root set of a polynomial about the imaginary axis to the approximation of some bases for the dominant eigenspaces of the matrices $(X_i \pm I_n)^h$ for a selected integer h and of the associated eigenvalues of C_p. We cover these tasks in the next two sections.

Algorithm 6.1. Splitting the root set of a polynomial about the imaginary axis.

1. *Perform iterations (6.1), by inverting matrices* X_i *by means of fast and numerically stable algorithms of [XXG12], [XXCB14], and [P15, Section 9.8].*
2. *When the iterations produce a matrix* X *with all eigenvalues lying near 1 and* -1*, compute two matrices* B_- *and* B_+ *whose columns approximate two bases for the dominant eigenspaces of the matrices* $(X_i \pm I_n)^h$ *for a selected integer* h*, respectively. The matrix* C_p *shares these eigenspaces by virtue of Theorem 12.2.*
3. *Compute and output approximations to the eigenvalues of the matrix* C_p *associated with those eigenspaces.*

Remark 6.1. *We have extended Theorem 3.1 to the companion matrix* C_p *and can extend it further to a matrix* $M = aC_p + cI_n$ *for complex numbers* $a \neq 0$ *and* c *if no eigenvalues of* M *lie on the imaginary axis. Then we can extend the algorithm to splitting the set of the eigenvalues into two sets separated by any straight line on the complex plane. Furthermore, unlike the polynomial map* $p(x)$ *to* $q(y) \to p(ay + c)$*, the map* $M \to aM + cI_n$ *is numerically stable and* enables numerically stable tiling *of Section 3. By performing an inversion of a shifted matrix* M*, we can extend Algorithm 6.1 to splitting root sets about a circle on the complex plane.*

7 Completion of Algorithm 6.1 by using Gaussian multipliers

7.1 A basic theorem for basis approximation

Hereafter a matrix is said to be an *approximate matrix basis* of a linear space if its columns approximate a basis for that space. Given an $n \times n$ matrix A that has dominant eigenspace of dimension r, we seek an approximate matrix basis B of that eigenspace. The following theorem directs us to the solution to this task given by the matrix $B = AH$ for an appropriate $n \times r$ matrix H.

Theorem 7.1. *Suppose that A is a diagonalizable $n \times n$ matrix with eigendecomposition $AV = V\Lambda$ of (12.1), $V = (V_2 \mid V_0)$ for an $n \times r$ block V_2 and an $n \times (n-r)$ block V_0,*

$$\Lambda = \operatorname{diag}(\lambda_i)_{i=1}^{n} = \operatorname{diag}(bI_r + E_2, E_0),$$

for $b = 2$ or $b = -2$, $||E_2|| + ||E_0|| \le \epsilon$, and a fixed small positive tolerance ϵ, so that the first r eigenvalues are close to b, and the $n-r$ remaining eigenvalues nearly vanish (e.g., this holds for the matrix $A = X_i \pm I_n$, an iterate X_i in (6.1), and a large integer i). Let

$$Y = V^{-1}H = \begin{pmatrix} Y_2 \\ Y_0 \end{pmatrix},$$

for an $n \times r$ matrix H, $Y_2 = (I_r \mid O_{n-r,r})Y$, and an $(n-r) \times r$ matrix Y_0. Then

$$AH = (bV_2Y_2 \mid O_{n-r,r}) + E \text{ and } ||E|| \le 2\epsilon\kappa(V)||H||. \tag{7.1}$$

Proof. Observe that
$AH = AVY = V\Lambda Y = V \operatorname{diag}(bI_r + E_2, E_0)Y$.
Therefore $AH = (bV_2Y_2 + V_2E_2Y_2 \mid V_0E_0Y_0)$, and so
$AH = (bV_2Y_2 \mid O_{n-r,r}) + E$ for $E = (V_2E_2Y_2 \mid V_0E_0Y_0)$.
It remains to estimate the norm $||E||$.
Note that $||E_2|| + ||E_0|| \le \epsilon$, $||(V_2 \mid V_0)|| \le ||V||$, and $||(Y_2 \mid Y_0)|| \le ||V||$ by assumptions. Therefore

$$\max\{||V_2E_2Y_2||, ||V_0E_0Y_0||\} \le \epsilon||V||\ ||Y||,$$

and so $||E|| \le 2\epsilon||V||\ ||Y||$.
Substitute $Y = V^{-1}H$ and deduce that
$||V||\ ||Y|| = ||V||\ ||V^{-1}H|| \le ||V||\ ||V^{-1}||\ ||H|| = \kappa(V)||H||$.
Combine the two latter bounds and obtain (7.1). □

Unless the matrix Y_2 is singular, bound (7.1) implies that, within the norm bound $||E|| \le 2\epsilon\kappa(V)||H||$ (that is, $||E|| \le 2\epsilon\kappa(V)$ if $||H|| = 1$), the matrix AH approximates a matrix basis of the eigenspace $\mathcal{V}$ generated by the columns of the matrix V_2.

In application at Stage 2 of Algorithm 6.1 we deal with the matrices $A = (X \pm I_n)^h$ and $B_\pm = AH$, respectively. We choose a matrix H, but do not know the matrix V. So we cannot say a priori how close to one another the column spans of the matrices AH and V_2 are, but we can readily check a posteriori whether $A^2H \approx bAH$ or not. Furthermore, by computing rank revealing factorization of [GL13, Section 5.4]) fro the matrix AH, we can verify that under perturbations within the norm bound $2\epsilon\kappa(V)||H||$ the matrix AH keeps full rank and is an approximate orthogonal matrix basis of the eigenspace $\mathcal{V}$. Computing the matrices AH is a little more costly for $A = (X_i \pm I_n)^h$ and $h \ne 0, 1$ than for $A = X_i \pm I_n$, but can improve convergence and the output accuracy significantly.

It remains to handle the following problems.

(i) The norm $||E||$ can be large if the condition number $\kappa(V)$ is large, that is, if the matrices C_p and $A = f(C_p)$ are close to derogatory matrices. We only have heuristic remedies of Example 5.1, but have never faced this problem in our tests.

(ii) Depending on the choice of a multiplier H, the product AH can be singular or ill-conditioned (that is, close to singular) even if the norm $||E||$ is small. We study this problem next.

7.2 Computation of an approximate matrix basis by using compression. Primal compression theorem

Consider an $n \times n$ matrix $A = A_r + E$ where $\text{rank}(A_r) = r < n$ and the norm $||E||$ is small. We are going to deal with multipliers H of reasonably bounded norms, and so $AH = A_rH + EH$ where the norm $||EH||$ is small. We will simplify our analysis by dropping the matrix E and assuming that the matrix $A = A_r$ has rank r, but we will allow more general multipliers H of size $n \times l$ for $r \le l = r + h < n$.

Theorem 7.2. The Primal Compression Theorem. *$\kappa(AH) \le \kappa(A)\kappa_{r,l}$, for an $n \times n$ matrix A of rank r, a Gaussian $n \times l$ matrix H, $r \le l = r + h < n$, and $\kappa_{r,l}$ of Section 13.*

Proof. Define compact SVD $A = S\Sigma T^H$ where S and T are $n \times r$ unitary matrices and Σ is a $r \times r$ diagonal matrix of the singular values of the matrix A.

Then $AH = S\Sigma G_{r,l}$ where $G_{r,l} = T^H H$ is a Gaussian $r \times l$ matrix by virtue of Lemma 13.1.

It follows that

$$||AH|| \le ||S||\ ||\Sigma||\ ||G_{r,l}||.$$

Substitute $||S|| = 1$, $||\Sigma|| = ||A||$, and $||G_{r,l}|| = \nu_{r,l}$ (cf. Section 13) and obtain

$$||AH|| \le ||A||\nu_{r,l}.$$

Moreover note that $(AH)^+AH = G^+_{r,l}G_{r,l} = \text{diag}(I_\rho, O_{l-\rho,l-\rho})$, for $\rho = \text{rank}(G_{r,l})$, and so

$$(AH)^+ = G^+_{r,l}\Sigma^{-1}S^H.$$

Therefore

$$||(AH)^+|| \le ||G^+_{r,l}|\ ||\Sigma^{-1}||\ ||S^H||.$$

Substitute $||S^H|| = 1$, $||\Sigma^{-1}|| = ||A^+||$, and $||G^+_{r,l}|| = \nu^+_{r,l}$ (cf. Section 13) and obtain

$$||(AH)^+|| \le ||A^+||\nu^+_{r,l}.$$

Hence

$$\kappa(AH) = ||AH||\ ||(AH)^+|| \le ||A||\nu_{r,l}||A^+||\nu^+_{r,l} = \kappa(A)\kappa_{r,l}. \tag{7.2}$$

□

Corollary 7.1. *Suppose that an $n \times n$ matrix A of rank r is well-conditioned and H is a Gaussian $n \times l$ matrix for $r \le l < n$. Then the column span of the matrix AH approximates a matrix basis of the dominant eigenspace of matrix A. Moreover, the matrix A is well-conditioned with a probability close to 1 and converging to 1 fast as $l - r$ grows.*

Proof. Combine Theorems 7.1 and 7.2 and Remark 13.1. □

The corollary shows some benefits of choosing a Gaussian $n \times l$ multiplier H where $l - r$ is positive and not close to 0 (empirically $l - r = 4$ is usually sufficient). By applying a rank-revealing factorization of the matrix AH, we can compress it into $n \times r$ an approximate matrix basis B of the dominant eigenspace of the matrix A. Such a compression is routine and noncostly if $l \ll n$ (see [GL13, Section 5.4] for rank-revealing factorization and see [HMT11, Propositions 10.2 and 10.3] for a similar policy for low-rank approximation of a matrix).

7.3 Extension to approximation of the eigenvalues of a companion matrix

Given an $n \times r$ approximate matrix basis B of rank $r < n$ for an eigenspace $\mathcal{U}$ (of dimension r) of a companion matrix C_p, the following algorithm approximates the associated eigenvalues of the matrix C_p. Theorem 12.1 implies its correctness.

Algorithm 7.1. *Extension to the approximation of the eigenvalues of a companion matrix.*

1. *Compute a* left inverse $B^{(I)}$ *of the matrix* B *(that is, a matrix such that* $B^{(I)}B = I_r$*).*

2. *Compute the matrix* $L = B^{(I)}C_pB$*.*

3. *Compute and output approximations to the eigenvalues of the matrix* L*.*

8 Completion of Algorithm 6.1 using structured multipliers

In our application to root-finding, the $n \times n$ matrix A of Corollary 7.1 has Toeplitz-like structure, the computation of the $n \times l$ product $B = AH$ involves order of $nl\log(n)$ flops, and the subsequent computation of the matrix $B^{(I)}C_pU$ requires order of nl^2 flops if H is a Gaussian $n \times l$ matrix. By choosing a Toeplitz-like $n \times l$ matrix H we can decrease both bounds below $cn\log(n)$, where the constant c is quite small if H is an $n \times r$ block of a circulant or f-circulant matrix of Section 14, and this constant decreases further if H belongs to the same Frobenius algebra $\mathcal{A}$ as the matrices A and C_p.

8.1 Computation of approximate bases: the Dual Compression Theorem

The following theorem shows that for $r \le l < n$ and a well-conditioned $n \times l$ multiplier H, the matrix AH is well-conditioned provided that $A = VG^T$, V and G are $n \times r$ matrices, V is a well-conditioned matrix of rank r and G is average $r \times n$ Gaussian matrix A, and we call such a matrix $A = VG^T$ an *average* $n \times n$ *matrix of rank* r under the Gaussian probability distribution or just an *average* $n \times n$ *Gaussian matrix of rank* r. The theorem holds for any well-conditioned $n \times l$ matrix H, and so we can endow it with Toeplitz-like structure; moreover we can choose it to be an $n \times l$ block of an f-circulant (e.g., circulant) matrix or a matrix in the associated Frobenius algebra $\mathcal{A}$.

Theorem 8.1. The Dual Compression Theorem. *Let* $A = VG^T$*, for a well-conditioned* $n \times r$ *matrix* V *of full rank* r*, a Gaussian* $n \times r$ *matrix* G*, a well-conditioned* $n \times l$ *matrix* H*,* $r \le l = r + h < n$*, and the number* $\kappa_{r,l}$ *of Section 13. Then* $\kappa(AH) \le \kappa(V)\kappa_{r,l}\kappa(H)$ *with probability 1.*

Proof. Let $V = S_V\Sigma_V T_V^H$ and $H = S_H\Sigma_H T_H^H$ be compact SVDs. Then

$$AH = VG^TH = S_V\Sigma_V T_V^H G^T S_H\Sigma_H T_H^H$$

where S_V, T_V^H, S_H, and T_H^H are unitary matrices of sizes $n \times r$, $r \times r$, $n \times l$, and $l \times l$, respectively, and where Σ_V and Σ_H are diagonal matrices of sizes $r \times r$ and $l \times l$, respectively. So

$$AH = S_V\Sigma_V G_{r,1}\Sigma_H T_H^H$$

where the $r \times l$ matrix

$$G_{r,1} = T_V^H G^T S_H$$

is Gaussian by virtue of Lemma 13.1.

Let $\Sigma_V G_{r,1}\Sigma_H = S_1\Sigma_1 T_1^H$ be compact SVD, where S_1 and T_1^H are unitary matrices of sizes $r \times r$ and $l \times l$, respectively, while Σ_1 is a diagonal matrix of size $l \times l$. Then

$$AH = S_V S_1\Sigma_1 T_1^H T_H^H.$$

Note that $S = S_V S_1$ and $T = T_1^H T_H^H$ are unitary matrices of sizes $n \times r$ and $l \times l$, respectively. Hence $AH = S\Sigma_1 T$ is compact SVD, and so

$$\kappa(AH) = \kappa(\Sigma_1) = \kappa(\Sigma_V G_{r,1}\Sigma_H).$$

As in the proof of equation (7.2) deduce that

$$\kappa(\Sigma_V G_{r,1}\Sigma_H) \le \kappa(V)\kappa(G_{r,1})\kappa(H),$$

substitute $\kappa(G_{r,1}) = \kappa_{r,1}$, and obtain the theorem. □

8.2 Extension from structured matrix basis to approximation of the eigenvalues

Suppose the column span of a well-conditioned Toeplitz-like $n \times l$ matrix $B_{n,l} = AH$, approximates a basis of an eigenspace $\mathcal{U}$ of dimension $r \le l$ for a companion matrix C_p. Then the algorithms of the previous section can approximate the associated eigenvalues but would destroy Toeplitz-like structure. Our alternative is to extend the heuristic algorithm of [GOSTZ10], which computes the so called CUR decomposition of a matrix and is highly efficient and dependable according to its extensive tests by the authors. It relies on the following result where $v_r(M) = \max_X |\det(X)|$ over all $r \times r$ submatrices X of a matrix M and is said to be the *maximal volume* of all such submatrics.

Theorem 8.2. *[GT01, Corollary 2.3]. Consider an $n \times l$ matrix $B = B_{n,l} = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix}$ with a nonsingular $r \times r$ leading block B_{11}. Write $|\det B_{11}|/v_r(B) = \nu$, $C = \begin{pmatrix} B_{11} \\ B_{21} \end{pmatrix}$, and $R = (B_{11} \mid B_{12})$ and let $||\cdot||_C$ denote the element-wise (Chebyshev) norm $||M|| \le n||M||_C$ for an $n \times n$ matrix M. Then $||B - CB_{11}^{-1}R||_C \le (r+1)\sigma_{r+1}(B))/\nu$.*

By virtue of the theorem, the rank-r matrix $CB_{11}^{-1}R$ approximates the matrix B within a factor of $(r+1)n/\nu$ from the optimal error bound $\sigma_{r+1}(B))$. ($CB_{11}^{-1}R$ is a CUR decomposition if the matrices B_{11} and $U = B_{11}^{-1}$ are unitary.)

The iterative algorithm of [GOSTZ10] converges fast to an $r \times r$ submatrix of the matrix $B = B_{n,l}$ having a reasonably large ratio $\nu < 1$. Such a nearly optimal rank-r approximation to the matrix B induces close approximation by the matrix CB_{11}^{-1} to its $n \times r$ matrix basis. The top $r \times r$ block of the matrix CB_{11}^{-1} is the identity matrix I_r, and so we immediately obtain the left inverse $(CB_{11}^{-1})^{(I)} = (I_r \mid O_{r,n-r})$.

The algorithm of [GOSTZ10] interchanges rows and columns of an input matrix, destroying Toeplitz-like structure, but we avoid this problem by means of the back and forth transition to Cauchy-like matrices [P90], [P15], [XXG12].

We plan to explore various alternative heuristic approaches. For example, for $l = r$, we can compute the $r \times n$ left inverse $B_{n,r}^{(I)} = B_{n,r}^T (B_{n,r}^T B_{n,r})^{-1}$. Clearly $B_{n,r}^{(I)} B_{n,r} = I_r$ because $B_{n,r} B_{n,r}^{(I)} B_{n,r} = B_{n,r}$. Is such computation numerically stable? The sharp bound $\kappa(M^T M) \le (\kappa(M))^2$ is acceptable if and only if M is sufficiently strongly well-conditioned.

9 Tiling by means of matrix computations

Suppose that the column sets of $n \times r_h$ matrices B_h approximate some bases for k eigenspaces $\mathcal{V}_h$ associated with k sets $\mathcal{S}_h$ of the eigenvalues of matrix C_p, for $h = 1, \dots, k$. The following algorithm computes an approximates matrix basis B for the eigenspace associated with the intersection $\cap_{h=1}^k \mathcal{S}_h$ of these k sets. For simplicity we let $k = 2$ because the extension to any integer k is immediate.

1. Compute $r_h \times (n - r_1 - r_2)$ matrices N_h, $h = 1, 2$, such that the columns of the $(r_1 + r_2) \times (n - r_1 - r_2)$ matrix $(N_1^T \mid N_2^T)^T$ form a basis for the null space of the matrix $(B_1 \mid B_2)$, so that $(B_1 \mid B_2)(N_1^T \mid N_2^T)^T = B_1 N_1^T + B_2 N_2^T = O_{n \times (n - r_1 - r_2)}$.
2. Compute and output the matrix $B = B_1 N_1$.

Correctness of the algorithm is readily verified.

Randomized algorithms of [PQ10] and [PZa] compute matrices N_1 and B by using $O((r_1 + r_2)^2 n)$ flops if B_1 and B_2 are general matrices, and we only need $O(n \log(n))$ flops if they are Toeplitz-like matrices.

10 Numerical tests

Numerical tests have been performed by Liang Zhao in the Graduate Center of the City University of New York, on a Dell server by using Windows system and MATLAB R2014a (with the IEEE

standard double precision), in particular the Matlab functions "cond()", "randn()", and "roots()". Polynomials of degree $n = 256$ of four types have been tested, namely, the polynomials of **Types I and II** had coefficients being i.i.d. standard Gaussian random real and complex variables, respectively, and benchmark polynomials of **Types III and IV** have been taken from the website of MPSolve (http://numpi.dm.unipi.it/mpsolve-2.2/bench.htm),
namely, $p(x) = x^n + 2x^{n-1} + ... + nx + (n+1)$ and $p(x) = x^{n+1} - 1$, respectively. In these initial tests our algorithms have been applied with standard Gaussian $n \times l$ random multipliers, i.e., matrices filled with i.i.d. standard Gaussian variables, for $l = r + 4$, and the eigenspaces of the matrices $X_i \pm I_n$ (that is, in the most primitive case where $h = 1$) have been approximated. Table 10.1 displays the number of iterations (6.1) performed and the error bound maximized for all roots and then taken average over 100 tests. The errors were computed in comparison with the data from the default MATLAB root-finding function "roots()".

Table 10.1: Root-finding by means of enhanced Cardinal method with oversampling

n	Type	iter	error
256	I	15.1	5.04e-14
256	II	13.4	1.77e-14
256	III	11.0	2.22e-14
256	IV	18.0	6.29e-12

Further tests for $n = 1024$ have shown very slow growth of the number iterations (6.1).

Acknowledgements: Our research has been supported by NSF Grant CCF–1116736 and PSC CUNY Awards 67699-00 45 and 68862–00 46.

References

[A66] L. Arnold, Uber die Nullstellenverteilung zufalliger Polynome (On the distribution of roots of random polynomials), *Math. Z*, **92** (1966), 12–18.

[B40] E.T. Bell, *The Development of Mathematics*, McGraw-Hill, New York, 1940.

[B68] C.A. Boyer, *A History of Mathematics*, Wiley, New York, 1968.

[B83] S. Barnett, *Polynomial and Linear Control Systems*, Marcel Dekker, New York, 1983.

[B85] J. R. Bunch, Stability of methods for solving Toeplitz systems of equations, *SIAM J. Sci. Stat. Comput.*, **6(2)**, 349–364, 1985.

[B96] D. A. Bini, Numerical Computation of Polynomial Zeros by Means of Aberth's Method, *Numerical Algorithms*, **13**, **3–4**, 179–200, 1996.

[BA80] R. R. Bitmead, B. D. O. Anderson, Asymptotically fast solution of Toeplitz and related systems of linear equations, *Linear Algebra and Its Applications*, **34**, 103–116, 1980.

[BB10] D.A. Bini, P. Boito, A fast algorithm for approximate polynomial GCD based on structured matrix computations. In: *Operator Theory: Advances and Applications*, **199**, pp. 155–173, Birkhäuser Verlag, Basel, 2010.

[BJ76] G. E. P. Box, G. M. Jenkins, *Time Series Analysis: Forecasting and Control*, Holden-Day, San Francisco, California, 1976.

[BP96] D. Bini, V. Y. Pan, Graeffe's, Chebyshev, and Cardinal's processes for splitting a polynomial into factors, *J. Complexity*, **12**, 492–511, 1996.

[BP00] C. Brunie, P. Picart, A fast version of the Schur–Cohn algorithm, *Journal of Complexity*, **16, 1**, 54–69, 2000.

[BSSY] Becker, R., Sagraloff, M., Sharma, V., Yap, C.: A simple near-optimal subsivision algorithm for complex root isolation based on Pellet's theorem and Newton iteration, arXiv: 1509.06231, December 2015.

[C96] J. P. Cardinal, On two iterative methods for approximating the roots of a polynomial, *Lectures in Applied Mathematics*, **32** (*Proceedings of AMS-SIAM Summer Seminar: Mathematics of Numerical Analysis: Real Number Algorithms* (J. Renegar, M. Shub, and S. Smale, editors), Park City, Utah, 1995), 165–188, American Mathematical Society, Providence, Rhode Island, 1996.

[CD05] Z. Chen, J. J. Dongarra, Condition Numbers of Gaussian Random Matrices, *SIAM. J. on Matrix Analysis and Applications*, **27**, 603–620, 2005.

[D88] J. Demmel, The Probability That a Numerical Analysis Problem Is Difficult, *Math. of Computation*, **50**, 449–480, 1988.

[DGS14] C. D'Andrea, A. Galligo, M. Sombra, Quantitative equidistribution for the solutions of systems of sparse polynomial equations, arXiv:1203.1843v3 [math.CV] 6 Aug 2014.

[DM89] C. J. Demeure, C. T. Mullis, The Euclid algorithm and the fast computation of cross–covariance and autocovariance sequences, *IEEE Trans. Acoust., Speech, Signal Processing*, **37**, 545–552, 1989.

[DM90] C. J. Demeure, C. T. Mullis, A Newton–Raphson method for moving-average spectral factorization using the Euclid algorithm, *IEEE Trans. Acoust., Speech, Signal Processing*, **38**, 1697–1709, 1990.

[DS01] K. R. Davidson, S. J. Szarek, Local Operator Theory, Random Matrices, and Banach Spaces, in *Handbook on the Geometry of Banach Spaces* (W. B. Johnson and J. Lindenstrauss editors), pages 317–368, North Holland, Amsterdam, 2001.

[E88] A. Edelman, Eigenvalues and Condition Numbers of Random Matrices, *SIAM J. on Matrix Analysis and Applications*, **9**, **4**, 543–560, 1988.

[ES05] A. Edelman, B. D. Sutton, Tails of Condition Number Distributions, *SIAM J. on Matrix Analysis and Applications*, **27**, **2**, 547–560, 2005.

[GL13] G. H. Golub, C. F. Van Loan, *Matrix Computations*, 4th edition, The Johns Hopkins University Press, Baltimore, Maryland, 2013.

[GOSTZ10] S. Goreinov, I. Oseledets, D. Savostyanov, E. Tyrtyshnikov, and N. Zamarashkin, How to find a good submatrix, in *Matrix Methods: Theory, Algorithms, Applications*, V. Olshevsky and E. Tyrtyshnikov, eds., World Scientific Publishing, pp. 247–256, 2010.

[GT01] S. A. Goreinov and E. E. Tyrtyshnikov, The maximal-volume concept in approximation by low-rank matrices, *Contemporary Mathematics*, **208**, 47–51, 2001.

[H59] A.S. Householder, Dandelin, Lobachevskii, or Graeffe, *Amer. Math. Monthly*, **66**, 464–466, 1959.

[H08] N. J. Higham, *Functions of Matrices: Theory and Computations*, SIAM, Philadelphia, 2008.

[HMT11] N. Halko, P. G. Martinsson, J. A. Tropp, Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Review*, **53, 2**, 217–288, 2011.

[HN08] C. P. Hughes, A. Nikeghbali, The zeros of random polynomials cluster uniformly near the unit circle, *Compos. Math.*, **144**, 734–746, 2008.

[JT70] M. A. Jenkins, J. F. Traub, A Three-Stage Variable-Shift Iteration for Polynomial Zeros and Its Relation to Generalized Rayleigh Iteration, *Numerische Math.*, **14**, 252–263, 1969/1970.

[JT72] M. A. Jenkins, J. F. Traub, A Three-Stage Algorithm for Real Polynomials Using Quadratic Iteration, *SIAM J. Numer. Anal.*, **7**, 545–566, 1970.

[K98] P. Kirrinnis, Partial fraction decomposition in $C(z)$ and simultaneous Newton iteration for factorization in $C[z]$. *Journal of Complexity*, **14**, 378–444, 1998.

[M80] M. Morf, Doubling algorithms for Toeplitz and related equations, *Proceedings of IEEE International Conference on ASSP*, 954–959, IEEE Press, Piscataway, New Jersey, 1980.

[MP13] J.M. McNamee, V.Y. Pan, *Numerical Methods for Roots of Polynomials*, Part 2 (XXII + 718 pages), Elsevier, 2013.

[P90] V. Y. Pan, On Computations with Dense Structured Matrices, *Math. of Computation*, **55, 191**, 179–190, 1990. Proceedings version in *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC'89)*, 34–42, ACM Press, New York, 1989.

[P95] V. Y. Pan, Optimal (up to polylog factors) sequential and parallel algorithms for approximating complex polynomial zeros, *Proc. 27th Ann. ACM Symp. on Theory of Computing (STOC'95)*, 741–750, ACM Press, New York, May 1995.

[P98/01] V. Y. Pan, Numerical computation of a polynomial GCD and extensions, *Information and Computation*, **167, 2**, 71–85, 2001, Proceedings version: Approximate polynomial Gcds, Padé approximation, polynomial zeros, and bipartite graphs, *Proc. 9th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'98)*, pages 68–77, ACM Press, New York, and SIAM Publications, Philadelphia, 1998.

[P00] V. Y. Pan, Approximating complex polynomial zeros: modified quadtree (Weyl's) construction and improved Newton's iteration, *J. of Complexity*, **16**, **1**, 213–264, 2000.

[P01] V. Y. Pan, *Structured Matrices and Polynomials: Unified Superfast Algorithms*, Birkhäuser/Springer, Boston/New York, 2001.

[P02] V. Y. Pan, Univariate polynomials: nearly optimal algorithms for factorization and rootfinding, *Journal of Symbolic Computation*, **33, 5**, 701–733, 2002. Proc. version in ISSAC'2001, pages 253–267, ACM Press, New York, 2001.

[P05] V. Y. Pan, Amended DSeSC power method for polynomial root-finding, *Computers and Math. (with Applications)*, **49**, **9**–**10**, 1515–1524, 2005.

[P15] V. Y. Pan, Transformations of matrix structures work again, *Linear Algebra and Its Applications*, **465**, 1–32, 2015.

[PQ10] V. Y. Pan, G. Qian, Randomized preprocessing of homogeneous linear systems, *Linear Algebra and Its Applications*, **432**, 3272–3318, 2010.

[PT13] Pan, V.Y., Tsigaridas, E.P.: On the Boolean complexity of the real root refinement. *Proc. Intern. Symposium on Symbolic and Algebraic Computation (ISSAC 2013)*, (M. Kauers editor), pp. 299–306, Boston, MA, June 2013. ACM Press, New York, 2013.

[PZa] V. Y. Pan, L. Zhao, Low-rank approximation of a matrix: novel insights, new progress, and extensions, arXiv 1510.06142[math.NA] (14 pages, 1 table), submitted on 21 Oct 2015, revised 15 Dec 2015.

[R87] J. Renegar, On the worst-case arithmetic complexity of approximating zeros of polynomials, *J. of Complexity*, **3, 2**, 90–113, 1987.

[S82] A. Schönhage, The fundamental theorem of algebra in terms of computational complexity, Report, Math. Department, Univ. Tübingen, Germany, 1982.

[S01] G. W. Stewart, *Matrix Algorithms, Vol II: Eigensystems*, SIAM, Philadelphia, 2001.

[SST06] A. Sankar, D. Spielman, S.-H. Teng, Smoothed Analysis of the Condition Numbers and Growth Factors of Matrices, *SIAM J. on Matrix Analysis and Applics.*, **28**, **2**, 446–476, 2006.

[VD94] P. M. Van Dooren, *Some Numerical Challenges in Control Theory*, *Linear Algebra for Control Theory, IMA Vol. Math. Appl.*, **62**, 1994.

[VdS70] A. Van der Sluis, Upper bounds on the roots of polynomials, *Numerische Math.*, **15**, 250–262, 1970.

[W69] G. T. Wilson, Factorization of the covariance generating function of a pure moving-average process, *SIAM Journal on Numerical Analysis*, **6**, 1–7, 1969.

[W07] D. S. Watkins, *The Matrix Eigenvalue Problem: GR and Krylov Subspace Methods*, SIAM, Philadelphia, PA, 2007.

[XXCB14] J. Xia, Y. Xi, S. Cauley, and V. Balakrishnan, Superfast and stable structured solvers for Toeplitz least squares via randomized sampling, *SIAM J. Matrix Anal. and Applications*, **35**, 44–72, 2014.

[XXG12] J. Xia, Y. Xi, M. Gu, A superfast structured solver for Toeplitz linear systems via randomized sampling, *SIAM J. Matrix Anal. Appl.*, **33**, 837–858, 2012.

Appendix

11 Some auxiliary results for polynomial computations

Some basic maps of the roots of a polynomial can be computed at linear or nearly linear arithmetic cost.

Theorem 11.1. *(*Root Inversion, Shift and Scaling*, cf. [P01].)*

(i) Given a polynomial $p(x)$ of (3.1) and two scalars a and b, one can compute the coefficients of the polynomial $q(x) = p(ax + b)$ by using $O(n \log(n))$ arithmetic operations. This bound decreases to $2n - 1$ multiplications if $b = 0$.

(ii) Reversing a polynomial inverts all its roots by involving no flops, that is,

$$p_{\rm rev}(x) = x^n p(1/x) = \sum_{i=0}^{n} p_i x^{n-i} = p_n \prod_{j=1}^{n} (1 - x x_j).$$

Note that by shifting and scaling the variable, we can move all roots of $p(x)$ into a fixed disc, e.g., $\{x : \ |x| \le 1\}$.

Theorem 11.2. *(*Dandelin's Root Squaring*, cf. [H59].)*

(i) Let a polynomial $p(x)$ of (3.1) be monic. Then $q(x) = (-1)^n p(\sqrt{x}\,) p(-\sqrt{x}\,) = \prod_{j=1}^{n} (x - x_j^2)$.

(ii) One can evaluate $p(x)$ at the k-th roots of unity, for $k > 2n$, and then interpolate to $q(x)$ by using $O(k \log(k))$ arithmetic operations overall.

Theorem 11.3. *(*The Cayley Maps.*)*

(i) The map $y = (x - a\sqrt{-1}\ )/(x + a\sqrt{-1}\ )$*, for any real nonzero scalar* a*, sends the real axis* $\{x :\ x \text{ is real}\}$ *onto the unit circle* $C(0,1) = \{y :\ |y| = 1\}$.

(ii) The converse map $x = a\sqrt{-1}\ (1-y)/(y+1)$ *sends the unit circle* $C(0,1)$ *onto the real axis.*

Theorem 11.4. *(See [VdS70].) Given a polynomial* $p(x)$ *of (3.1), write* $r_1 = \max_{j=1}^{n} |x_j|$ *and* $\gamma^+ = \max_{i=1}^{n} |p_{n-i}/p_n|$*. Then* $\gamma^+/n \le r_1 \le 2\gamma^+$.

12 Definitions and basic facts for eigenvalue computation

An $m \times n$ matrix M has an $n \times m$ *left inverse* matrix $X = M^{(I)}$ such that $XM = I_n$ if and only if it has full column rank n. In this case the Moore–Penrose pseudo inverse M^+ is a left inverse (cf. [GL13], [S01]). The left inverse is unique if and only if M is a nonsingular matrix, in which case $m = n$ and $M^{(I)} = M^{-1}$.

Definition 12.1. Eigenvalues, eigenvectors, eigenspaces.

- *A scalar* x *is an* eigenvalue *of a matrix* M *associated with an* eigenvector $\mathbf{v}$ *if* $M\mathbf{v} = x\mathbf{v}$.
- *The eigenvectors associated with an eigenvalue* x *or, more generally, with any set of the eigenvalues* $\mathcal{X}$ *of a matrix* M *form the* eigenspaces $\mathcal{S}(M,x)$ *and* $\mathcal{S}(M,\mathcal{X})$*, respectively, associated with the eigenvalue* x *and the set* $\mathcal{X}$ *of eigenvalues, respectively. A linear subspace* $\mathcal{S}$ *of* $\mathbb{C}^{n\times n}$ *is an eigenspace of a matrix* M *if and only if* $M\mathcal{S} = \{M\mathbf{v} : \mathbf{v} \in \mathcal{S}\} \subseteq \mathcal{S}$ *(see [S01, Definition 4.1.1]).*
- *An eigenvalue* x *of a matrix* M *is a root of the characteristic polynomial* $\det(xI - M)$*. The multiplicity of this root is the* algebraic multiplicity *of the eigenvalue* x*, denoted* $am(x)$*. The dimension* $gm(x) = \dim(\mathcal{S}(M,x))$ *is the* geometric multiplicity *of* x*, never exceeding* $am(x)$*. An eigenvalue* x *is* simple *if* $gm(x) = 1$.

For every $n \times n$ matrix A having eigenvalues λ_i with $gm(\lambda_i) = d_i$, $i = 1, \dots, h$, $\sum_{i=1}^{k} d_i = k$, where k is maximal, $k \le n$, there exists eigendecomposition

$$AV = V\Lambda \tag{12.1}$$

where $\Lambda = \mathrm{diag}(\lambda_i)_{i=1}^{k}$ is the diagonal matrix of the eigenvalues of the matrix A and V is an $n \times k$ matrix whose columns are its right eigenvectors, associated with these eigenvalues. The matrix A is called *diagonalizable* if $k = n$ and is called *derogatory* otherwise. In the latter case infinitesimal perturbations of the matrix can turn the matrix into diagonalizable, thus changing its eigenvalues and eigenvectors dramatically. If a matrix is close to derogatory, then the computation of its eigenvalues and eigenvectors is numerically unstable and requires high precision of computing.

Next, given a basis for an eigenspace of a matrix, let us approximate the associated eigenvalues.

At first suppose that an eigenvalue μ of a matrix M is simple and an associated eigenvector $\mathbf{v}$ (unique up to scaling) generates the associated eigenspace. Clearly, in this case the eigenvalue μ is given by the *Rayleigh quotient* $\mu = \frac{\mathbf{v}^T M\mathbf{v}}{\mathbf{v}^T\mathbf{v}}$ as well as by a *simple quotient* $f(\mu) = \frac{(M\mathbf{v})_i}{(\mathbf{v})_i}$ for any integer i, $1 \le i \le n$, such that $(\mathbf{v})_i \ne 0$. Here $(M\mathbf{v})_i$ and $(\mathbf{v})_i$ denote the ith coordinates of the vectors $M\mathbf{v}$ and $\mathbf{v}$, respectively.

The following simple theorem naturally extends the above expressions to the case where we are given a basis for any eigenspace of a matrix.

Theorem 12.1. *(*Decreasing the Eigenproblem Size to the Dimension of an Eigenspace*, cf. [W07, Section 2.1].)*

Suppose that an $n \times r$ *matrix* B *has rank* $r \le n$*, has a left inverse* $B^{(I)}$*, and its columns generate a linear space* $\mathcal{U}$*. Then*

(i) $\mathcal{U}$ *is an eigenspace of* M *if and only if there exists a matrix* $L \in \mathbb{C}^{r\times r}$ *such that* $L = B^{(I)}MB$ *(that is,* $MB = BL$*),*

(ii) every eigenvalue of the matrix L is an eigenvalue of the matrix M,

(iii) $MB\mathbf{v} = xB\mathbf{v}$ if $L\mathbf{v} = x\mathbf{v}$,

(iv) the matrix L is unique, that is, its choice is independent of the choice of a matrix U and its left inverse $U^{(I)}$, and so $L = B^H MB$, for a unitary matrix B.

Remark 12.1. *If the dimension r is small relative to n, then the arithmetic as well as Boolean cost of the approximation of the eigenvalues of an $r \times r$ matrix L is small. Even if the dimension r is large but the matrix L is structured and readily invertible, we can approximate its eigenvalues at a lower cost by applying Rayleigh quotient iteration (enhancing the Inverse Power method). If the matrix L is unstructured and the dimension r is large, we can try to decrease it by using tiling described at the end of Section 3.*

Theorem 12.2. *(*The Eigenproblems for a Matrix and Its Function.*)*

Suppose that M is a square matrix and that a rational function $f(x)$ is defined on its spectrum.

(i) Then $f(M)\mathbf{v} = f(x)\mathbf{v}$ if $M\mathbf{v} = x\mathbf{v}$.

(ii) Let $\mathcal{U} = \mathcal{U}_{\mu,f}$ denote the eigenspace of the matrix $f(M)$ associated with its eigenvalue μ. Then this is an eigenspace of the matrix M associated with all its eigenvalues x such that $f(x) = \mu$.

(iii) The space $\mathcal{U}$ has dimension 1 and is associated with a single eigenvalue of M if μ is a simple eigenvalue of $f(M)$.

Proof. We readily verify part (i), which implies parts (ii) and (iii). See also [H08, Theorem 1.13]. □

13 Random matrices

A matrix H is *Gaussian* if its entries are *i.i.d.* Gaussian variables.

Lemma 13.1. *(*Rotational invariance of a Gaussian matrix.*) Suppose that k, m, and n are three positive integers, G is an $m \times n$ Gaussian matrix, and S and T are $k \times m$ and $n \times k$ orthogonal matrices, respectively. Then SG and GT are Gaussian matrices.*

For the sake of brevity and simplicity we state the following estimates for real matrices, but similar estimates in the case of complex matrices can be found in [D88], [E88], [CD05], and [ES05].

Write $\nu_{m,n} = ||G||$, $\nu^+_{m,n} = ||G^+||$, and $\kappa_{m,n} = \kappa(G)$, for a Gaussian $m \times n$ matrix G, and write $\mathbb{E}(v)$ for the expected value of a random variable v.

Theorem 13.1. *(Cf. [DS01, Theorem II.7].) Suppose that m and n are positive integers, $h = \max\{m, n\}$, $t \geq 0$, and $z \geq 2\sqrt{h}$. Then*

Probability$\{\nu_{m,n} > z\} \leq \exp(-(z - 2\sqrt{h})^2/2)$ *and*

Probability$\{\nu_{m,n} > t + \sqrt{m} + \sqrt{n}\} \leq \exp(-t^2/2)$.

Theorem 13.2. *Let $\Gamma(x) = \int_0^\infty \exp(-t)t^{x-1}dt$ denote the Gamma function and let $x > 0$. Then*

(i) Probability $\{\nu^+_{m,n} \geq m/x^2\} < \frac{x^{m-n+1}}{\Gamma(m-n+2)}$ *for $m \geq n \geq 2$,*

(ii) Probability $\{\nu^+_{n,n} \geq x\} \leq 2.35\sqrt{n}/x$ *for $n \geq 2$,*

(iii) $\mathbb{E}(\nu^2_{F,m,n}) = m/|m - n - 1|$, provided that $|m - n| > 1$, and

(iv) $\mathbb{E}(\nu^+_{m,n}) \leq e\sqrt{m}/|m - n|$, provided that $m \neq n$.

Proof. See [CD05, Proof of Lemma 4.1] for part (i), [SST06, Theorem 3.3] for part (ii), and [HMT11, Proposition 10.2] for parts (iii) and (iv). □

Remark 13.1. *Theorem 13.2 provides probabilistic upper bounds on $\nu^+_{m,n}$. They are reasonable already for square matrices, for which $m = n$, but become much stronger as the difference $|m - n|$ grows large. Theorems 13.1 and 13.2 combined imply that an $m \times n$ Gaussian matrix is very well-conditioned if the integer $m - n$ is at least moderately large and that such a matrix can still be considered well-conditioned (possibly with a grain of salt) if the integer $|m - n|$ is small or even vanishes. These properties are immediately extended to all submatrices because they are also Gaussian.*

14 Computations with circulant and Toeplitz-like matrices

An $m \times n$ *Toeplitz* matrix $T = (t_{i-j})_{i,j=0}^{m-1,n-1}$ is defined by its first column and first row. The special case is an $n \times n$ *circulant* matrix $Z_1(\mathbf{v}) = (v_{i-j \mod n})_{i,j=0}^{n-1}$, defined by its first column vector $\mathbf{v} = (v_i)_{i=0}^{n-1}$.

Multiply all superdiagonal entries of $Z_1(\mathbf{v})$ by a fixed complex number $f \neq 0$ and obtain an *f-circulant* matrix $Z_f(\mathbf{v})$.

Let $J = J_n$ denote the $n \times n$ *reversion matrix*, filled with 0s except for the anti-diagonal filled with 1s, $J(v_i)_{i=1}^n = (v_{n+1-i})_{i=1}^n$ for any vector $(v_i)_{i=1}^n$.

For two distinct complex numbers e and f, an $n \times n$ matrix $M = \frac{1}{e-f} \sum_{i=1}^d Z_e(\mathbf{g}_i) Z_f(J\mathbf{h}_i)$ can be defined by a generator matrices $G = (g_i)_{i=1}^d$ and $H = (h_i)_{i=1}^d$. If d is small in context, then M is called a *Toeplitz-like* matrix, defined by its *displacement generator* (G, H) of length d. The minimum length for a fixed matrix M is called its *displacement rank*, denoted $d(M)$.

$n \times k$ and $k \times n$ blocks of an $n \times n$ circulant or f-circulant matrices are said to be *subcirculant* and *f-subcirculant* matrices, respectively. We may readily define an $m \times n$ Toeplitz-like matrix M by allowing to use e-subcirculant and f-subcirculant matrices in the above expression.

Theorem 14.1. *(i) If for two matrices M and N we can define their sum, product or inverses, then $d(M+N) \leq d(M) + d(N)$, $d(MN) \leq d(M) + d(N)$, and $d(M^{-1}) = d(M)$. Furthermore $d(cM) = d(M)$ for any scalar $c \neq 0$. (ii) If the matrices M and N are represented by their displacement generators of length $l(M)$ and $l(N)$ and if there exist their sum, product or inverses, then we need no flops in order to obtain a displacement generator of length $d(M) + d(N)$ for the matrix $M + N$; it is sufficient to use $O(d(M)d(N)n\log(n)$ flops in order to compute a displacement generator of length $d(M) + d(N)$ for the matrix MN; it is sufficient to use $(d(M)^2 n \log^2(n))$ flops in order to compute a generator of length $d(M)$ for the inverse matrix M^{-1} and to compute a vector $\mathbf{x} = M^{-1}\mathbf{f}$ satisfying a nonsingular linear system of n equations $M\mathbf{x} = \mathbf{f}$. (iii) Furthermore we can approximate the vector $\mathbf{x} = M^{-1}\mathbf{f}$ by using $O(d(M)^2 n \log^2(n))$ flops.*

Proof. The algorithms [XXG12] and [XXCB14] accelerated in [P15, Section 9.8] by a factor of $\log(n)$ support approximation of the vector $\mathbf{x} = M^{-1}\mathbf{f}$. The proofs of the other claimed estimates can be found in [P01]. □